

REMARKS/ARGUMENTS

Claim 1 has been amended to move the recitation “at least some of the frames contain multiple data packets” from the preamble of the claim to the body of the claim. Claim 1 has also been amended to explicitly recite “delaying transmission of some of the data packets so that at least some of the frames each contain multiple data packets” and in particular, step (b) has been amended to recite “(b) upon delaying packet transmission for the certain time interval, transmitting each frame in a second set of the frames which are not filled with at least some of the data packets so that said each frame in the second set of the frames cannot contain an additional data packet in order to ensure that said each data packet is transmitted in at least one of the frames no later than the certain time interval after the respective time of said each data packet in the time sequence.” Dependent claims 13, 25, and 35 also have been amended to recite “delaying transmission of some of the I/O request data packets by a certain time interval so that at least some of the network transmission frames each contain multiple I/O request data packets, ...” Support for this packet transmission delay so that at least some of the frames each contain multiple data packets is found in applicant’s FIG. 3 in comparison to FIG. 2 (delay of no more than “x” in FIG. 3 for filling MTU packet 55 with multiple I/O request packets 41, 42, 43; and delay of “x” for transmitting I/O request packet 44 in MTU frame 56), steps 92, 93, and 94 in applicant’s FIG. 6, and steps 162, 163, 164, 165, 166, and 167 in applicant’s FIG. 15, as described in applicant’s specification paragraph [00030] on page 8 line 22 to page 9 line 10; page 14 lines 1-8; page 20 line 15 to page 21 line 3; and page 26 lines 9-15.

On page 2 of the Official Action, claim 1 was rejected under 35 U.S.C. 103(a) as being unpatentable over Karpoff (U.S. Pat. 7,299,290) in view of Gunaseelan et al. (U.S. Patent Application Pub. 2002/0097750). In reply, claim 1 has been amended, as discussed above, to more clearly distinguish Karpoff and Gunaseelan. It is respectfully submitted that neither Karpoff nor Gunaseelan disclose “delaying transmission of some of the data packets so that at least some of the frames each contain multiple data packets,” nor does Karpoff nor Gunaseelan alone or in combination suggest the applicant’s method of delaying transmission of some of the data packets so that at least some of the frames each contain multiple data packets by the combination of steps (a) and (b) as recited in claim 1 as amended.

For example, page 2 of the Official Action cites Karpoff col. 12, lines 7-10. Karpoff col. 12, lines 5-10 say:

Towards this end, each computer attached to the Ethernet LAN, which includes both the clients 10) and server 12 (e.g., with reference to FIG. 7, transmits information in a way that accommodates a predefined packet size called a Maximum Transmission Unit (MTU), the value of which is 1,518 bytes for Ethernet. These packets are addressed at the LAN level using an Ethernet header that contains a unique MAC address. Every computer that attaches to an Ethernet network has a unique 48-bit MAC address, distinct from the 32-bit IP address, to determine which node on the LAN is the correct destination for the packet.

The disclosure of “transmit[ting] information in a way that accommodates a predefined packet size called a Maximum Transmission Unit (MTU)” is not sufficiently detailed to disclose or suggest the particular way in which the method of applicant’s claim 1 as amended

successively joins data packets from the time sequence into the frames and delays transmission of some of the data packets so that at least some of the frames each contain multiple data packets.

Page 3 of the Official Action cites Gunaseelan et al. paragraph [0026] lines 16-18 and 21-22 for disclosing packet transmission in frames no later than a certain time interval in a time sequence. Gunaseelan et al. paragraph [0026] lines 16-22 say:

As illustrated in the FIG. 4, such an on time delivery is shown where packet P1 is delivered at time T1, packet P2 is delivered at time t2, and so on. In addition to delivering packets of uniform size at regular intervals, the delivery system 200 also supports delivery of variable sized packets at fixed intervals in time (FIG. 5) such as when delivering I, P, or B frame data every $1/30^{\text{th}}$ of a second.

In claim 1 as amended, applicant's "certain time interval" is a limit of delay time of packet transmission, and this limit of delay time of packet transmission is reached for a packet in a frame that is transmitted when the frame is not completely filled with data packets. In simplified terms, a frame is transmitted in step (a) upon filling a frame with packets for a frame that cannot contain an additional packet, or a frame is transmitted in step (b) upon a packet transmission delay of a certain time interval for a frame that has not yet become filled with packets. It is respectfully submitted that the use of a certain time interval for control of packet transmission from frames as recited in applicant's claim 1 as amended is not suggested by Gunaseelan's delivery of the packets without missing deadlines.

On page 3 of the Official Action, claims 2-4 and 11 were rejected under 35 U.S.C. 103(a) as being unpatentable over Karpoff (U.S. Pat. 7,299,290) in view of Gunaseelan et al. (U.S. Patent Application Pub. 2002/0097750) further in view of Kouloheris et al. (U.S. Pat. 5,915,094). Applicant respectfully traverses in view of the amendment of the base claim 1. It is respectfully submitted that Kouloheris et al. does not disclose or suggest the elements of the base claim 1 that are missing from the proposed combination of Karpoff and Gunaseelan.

With respect to applicant's dependent claim 2, pages 4-5 of the Official Action cite Kouloheris et al. col. 10 line 30 for disclosing a timer interrupt routine 280. Kouloheris et al. col. 10 line 30 says: " The interrupts from the timer 280 in the stream controller 20 are used to define fixed interval network and disk cycles. ... The two real time processes, the disk read processes 460 and the network transmission processes 420 are initiated by timer interrupts which start the respective cycles." It is respectfully submitted that this general use a timer interrupt to start a network transmission process does not suggest that Karpoff and Kouloheris should be combined and modified to use the applicant's particular main routine and a timer interrupt routine as recited in claim 2 to perform the method of applicant's claim 1 as amended.

With respect to applicant's dependent claim 3, page 5 of the Official Action cites Kouloheris et al. col. 9 lines 28-31 as pertinent to separately joining read I/O requests together for transmission, and separately joining write I/O requests together for transmission. Kouloheris et al. col. 9 lines 28-31 say: "In each disk cycle the disk read process issues a read command for each I/O request to retrieve a portion of the data unit requested by the I/O request, thus dividing each I/O request into a sequence of read commands." It is respectfully submitted that this

general teaching of dividing a I/O request into a sequence of disk read commands does not suggest the applicant's method of separately joining read I/O requests together for transmission, and separately joining write I/O requests together for transmission, so that the I/O request data packets have an ordering in the frames that is different from the ordering of the I/O request data packets in the time sequence, as claimed in applicant's claim 3.

With respect to applicant's dependent claim 4, page 5 of the Official Action cites Kouloheris et al. col. 17 line 3 as pertinent to some of the read I/O request packets being moved in front of some of the write I/O request packets in some of the frames. Kouloheris et al. col. 17 line 3 says: " In addition to the four VCR interface commands already described so far, READ and WRITE commands allow the host to load data into the disk array 60, and read the data back to verify its integrity." It is not understood how this general disclosure of a host using READ and WRITE commands for loading data into a disk array and reading the data back to verify its integrity suggests the applicant's specific method in which some of the read I/O request data packets are moved in front of some of the write I/O request data packets in some of the frames, as recited in applicant's claim 4. For example, Kouloheris should not be reading back data to verify its integrity before the data has been written to the disk array.

. With respect to applicant's dependent claim 11, page 5 of the Official Action cites Kouloheris et al. col. 14 lines 23-24 as pertinent to measuring loading on the data network, and dynamically adjusting the duration of the certain time interval based on the measured loading of the data network, the duration of the certain time interval being increased for increased loading on the data network. Kouloheris et al. col. 14, lines 23-24 say: "Having large stripe groups does increase the delay in starting a new stream in a very heavily loaded system, but at the same

time reduces the probability of rejecting the request for starting a new stream due to hot spots appearing on the stripe group being accessed by the new request.” In reply, it is respectfully submitted that the amendment of claim 1 has more particularly defined the meaning and use of “the certain time interval” so that Kouloheris is no longer pertinent. In other words, the fact that large stripe groups in a RAID system may increase the delay in starting a new stream in a heavily loaded system does not suggest that “the certain time interval” of applicant’s claim 1 should be dynamically adjusted based on measured loading of the data network so that the duration of the certain time interval is increased for increased loading on the data network, as recited in applicant’s claim 11.

On page 6 of the Official Action, claims 5-10 were rejected under 35 U.S.C. 103(a) as being unpatentable over Karpoff (U.S. Pat. 7,299,290) in view of Gunaseelan et al. (U.S. Patent Application Pub. 2002/0097750) further in view of Madukkarumukumana et al. (U.S. Patent Application Pub. 2005/0030972). Applicant respectfully traverses in view of the amendment of the base claim 1. It is respectfully submitted that Madukkarumukumana does not disclose or suggest the elements of the base claim 1 that are missing from the proposed combination of Karpoff and Gunaseelan.

With respect to applicant’s dependent claim 7 (and dependent claims 8-10 which depend on claim 7), page 6 of the Official Action cites Madukkarumukumana page 3 paragraph [0028] as pertinent to data packets are stored in a range of addresses of memory, a certain number of frames are preallcoated in another region of memory, and the data packets are joined by transfer

of the data packets from the range of addresses in memory to the preallocated frames in memory.

However, Madukkarumukumana page 3 paragraph [0028] says:

[0028] Dividing an anonymous buffer 214a . . . 214n into packet regions 216a . . . 216n also allows the protocol processing application 120 to use a packet region to save the data payload in a full sized incoming packet, and link the packet region to the appropriate protocol control block (PCB) context of the session the packet belongs to. A packet region 216a . . . 216n is a region of memory in the anonymous buffers 214a . . . 214n, where the size of a packet region 216a . . . 216n may be equal to the size of the underlying transport's MTU, and a full sized incoming packet may be equal to the size of the underlying transport's MTU. The packet regions may be recycled, i.e., reused, after packets have been reassembled

In other words, Madukkarumukumana paragraph [0028] discloses dividing an anonymous buffer into packet regions in which the size of a packet region may be equal to the size of the underlying transport's MTU. It is respectfully submitted that Madukkarumukumana does not suggest joining packets in a MTU or transferring a data packet from a packet region to a MTU region, because in Madukkarumukumana the packet region is the same as the MTU region. Moreover, if the size of a data packet is equal to the size of the underlying transport's MTU, then one would not be "joining" packets in the MTU frame. Instead, one would put each packet in a corresponding MTU frame and transmit the frame without delaying transmission of the packet once the packet is put in its corresponding MTU frame.

On page 7 of the Official Action, claims 12, 15-20, 24, 27-34 and 37-41 were rejected under 35 U.S.C. 103(a) as being unpatentable over Karpoff (U.S. Pat. 7,299,290) in view of Kouloheris et al. (U.S. Pat. 5,915,094). Applicant respectfully traverses.

Page 8 of the Official Action cites Karpoff col. 12 lines 5-10 as pertinent to “the host processor joining data packets from different ones of the on-line transaction processing applications in the same network transmission frames to more completely fill the network transmission frames.” Karpoff col. 12 lines 5-10 say:

Towards this end, each computer attached to the Ethernet LAN, which includes both the clients 10) and server 12 (e.g., with reference to FIG. 7, transmits information in a way that accommodates a predefined packet size called a Maximum Transmission Unit (MTU), the value of which is 1,518 bytes for Ethernet. These packets are addressed at the LAN level using an Ethernet header that contains a unique MAC address. Every computer that attaches to an Ethernet network has a unique 48-bit MAC address, distinct from the 32-bit IP address, to determine which node on the LAN is the correct destination for the packet.

The disclosure of “transmit[ting] information in a way that accommodates a predefined packet size called a Maximum Transmission Unit (MTU)” is not sufficiently detailed to disclose or suggest “the host processor joining data packets from different ones of the on-line transaction processing applications in the same network transmission frames to more completely fill the network transmission frames” as recited in applicant’s claim 12. Nor does Kouloheris et al. disclose or suggest “the host processor joining data packets from different ones of the on-line

transaction processing applications in the same network transmission frames to more completely fill the network transmission frames” as recited in applicant’s claim 12.

When determining whether a claim is obvious, an examiner must make “a searching comparison of the claimed invention – *including all its limitations* – with the teaching of the prior art.” *In re Ochiai*, 71 F.3d 1565, 1572 (Fed. Cir. 1995) (emphasis added). Thus, “obviousness requires a suggestion of all limitations in a claim.” *CFMT, Inc. v. Yieldup Intern. Corp.*, 349 F.3d 1333, 1342 (Fed. Cir. 2003) (*citing In re Royka*, 490 F.2d 981, 985 (CCPA 1974)). Moreover, as the Supreme Court recently stated, “*there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness.*” *KSR Int’l v. Teleflex Inc.*, 127 S. Ct. 1727, 1741 (2007) (quoting *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006) (emphasis added)).

The problem that the inventor is trying to solve must be considered in determining whether or not the invention would have been obvious. The invention as a whole embraces the structure, properties and problems it solves. *In re Wright*, 848 F.2d 1216, 1219, 6 U.S.P.Q.2d 1959, 1961 (Fed. Cir. 988). As described in applicant’s specification on page 2 line 19 to page 3 line 6 and on page 8 lines 1-21, the applicant discovered a problem that in an on-line transaction processing system employing network storage, there is a significant degradation in performance under high loading conditions due to inefficient packing of block-level I/O requests into the MTU frames. Applicant respectfully submits that discovery of this problem was unexpected and surprising. The Official Action does not recognize the existence of this problem and therefore is itself evidence that discovery of the problem was unexpected. Nevertheless, it should be understood that the problem

does in fact arise when a conventional network transmission protocol stack is invoked by different on-line transaction processing applications executed by a host processor, because invocation of the conventional network transmission protocol stack by one on-line transaction processing application causes transmission of an MTU frame over the data network independent of invocation of the conventional network transmission protocol stack by another on-line processing application.

Applicant also respectfully submits that the discovery of this problem and its solution was not obvious because the problem must have existed for many years prior to its discovery and solution by applicant. For example, the IP protocol stacks in network storage systems have been known at least since the Dec. 6, 1994 priority date of Kouloheris, as evidenced by the disclosure in Kouloheris. See, for example, Kouloheris FIGS. 1-3 and Col. 20 line 10 et seq. See Fromson v. Advance Offset Plate, Inc., 755 F.2d 1549, 1557, 225 U.S.P.Q. 26, 32-33 (Fed. Cir. 1985) (It is at best bizarre to assert that the subject matter claimed was merely an obvious extension of technology when none skilled in the art attempted such “extension” during the seven years when alleged economic advantages of such technology were available).

With respect to applicant’s dependent claim 16, page 8 of the Official Action cites Kouloheris col. 17 line 3. Kouloheris col. 17 line 3 says: “In addition to the four VCR interface commands already described so far, READ and WRITE commands allow the host to load data into the disk array 60, and read the data back to verify its integrity.” It is not understood how this general disclosure of a host using READ and WRITE commands for loading data into a disk array and reading the data back to verify its integrity suggests the method of applicant’s claim 16 of separately joining the read I/O request data packets together for transmission to the network

block storage, and separately joining the write I/O request data packets together for transmission to the network block storage.

With respect to applicant's dependent claim 17, page 8 of the Official Action again cites Kouloheris col. 17 line 3. It is not understood how this general disclosure of a host using READ and WRITE commands for loading data into a disk array and reading the data back to verify its integrity suggests the method of applicant's claim 17 of moving some of the read I/O request data packets in front of some of the write I/O request data packets in some of the frames.

With respect to applicant's dependent claim 18, page 8 of the Official Action again cites Kouloheris col. 17 line 3. It is not understood how this general disclosure of a host using READ and WRITE commands for loading data into a disk array and reading the data back to verify its integrity suggests the method of applicant's claim 18 of turning on and off the joining of the I/O request data packets.

With respect to applicant's dependent claim 19, page 8 of the Official Action again cites Kouloheris col. 17 line 3. It is not understood how this general disclosure of a host using READ and WRITE commands for loading data into a disk array and reading the data back to verify its integrity suggests the method of applicant's claim 19 of turning on and off the joining of the I/O request data packets in which the joining is turned off during a bulk transfer of database data.

With respect to applicant's independent claim 24 and dependent claims 32-33, applicant respectfully traverses the rejection for the reasons given above with respect to the rejection of applicant's claim 12. Moreover, with respect to applicant's dependent claim 33, it is not seen where Karpoff or Kouloheris discloses or suggests "re-programming the network attached

storage to bunch I/O replies into frames for transmission from the network attached storage over the data network to the host processor.”

With respect to applicant’s dependent claim 28, applicant further traverses for the reasons given above with respect to applicant’s claim 16.

With respect to applicant’s dependent claim 29, applicant further traverses for the reasons given above with respect to applicant’s claim 17.

With respect to applicant’s dependent claim 30, applicant further traverses for the reasons given above with respect to applicant’s claim 18.

With respect to applicant’s independent claim 34, applicant respectfully traverses for the reasons given above with respect to applicant’s claim 12.

With respect to applicant’s dependent claim 38, applicant further traverses for the reasons given above with respect to applicant’s claim 16.

With respect to applicant’s dependent claim 39, applicant further traverses for the reasons given above with respect to applicant’s claim 17.

With respect to applicant’s dependent claim 40, applicant further traverses for the reasons given above with respect to applicant’s claim 18.

On page 12 of the Official Action, applicant’s dependent claims 13-14, 25-26 and 35-36 were rejected under 35 U.S.C. 103(a) as being unpatentable over Karpoff (U.S. Pat. 7,299,290) in view of Kouloheris et al. (U.S. Patent 5,915,094) further in view of Gunaseelan et al. (U.S. Patent Application Pub. 2002/0097750). Applicant respectfully traverses. Neither Gunaseelan nor Kouloheris appear to disclose or suggest the elements of the base claims 12, 24, and 34 that

are missing from Karpoff. In addition, claims 13, 25, and 35 have been amended to recite “delaying transmission of some of the I/O request data packets by a certain time interval so that at least some of the network transmission frames each contain multiple I/O request data packets,” so as to further distinguish Karpoff and Gunaseelan, for the reasons discussed above with respect to applicant’s claim 1.

With respect to applicant’s dependent claims 14, 26, and 36, applicant further traverses for the reasons given above with respect to applicant’s claim 11.

On page 13 of the Official Action, applicant’s dependent claims 21-23 and 42-43 were rejected under 35 U.S.C. 103(a) as being unpatentable over Karpoff (U.S. Pat. 7,299,290) in view of Kouloheris et al. (U.S. Patent 5,915,094) further in view of Madukkarumukumana et al. (U.S. Patent Application Pub. US 2005/0030972). Applicant respectfully traverses. Madukkarumukumana et al. does not appear to disclose or suggest the elements of the base claims 12 and 34 that are missing from Karpoff and Kouloheris. In addition, with respect to the limitations added by claims 21 and 42 (and incorporated by reference into dependent claims 22-23 and 42-43), applicant respectfully submits that Madukkarumukumana et al. does not disclose or suggest “joining the data packets during transfer of the data packets from the range of addresses in memory to the preallocated frames in memory” for the reasons given above with respect to the rejection of applicant’s claim 7.

With respect to the further limitations added in dependent claims 23, page 14 of the Official Action cites Madukkarumukumana figure 1, page 1, paragraph [0015], lines 1-17, and paragraph [0016]. Lines 17-20. However, it is not understood how Madukkarumukumana figure

1, page 1, paragraph [0015], lines 1-17, or paragraph [0016]. Lines 17-20 discloses or suggests ““bunching I/O replies into frames for transmission from the network attached storage over the data network to the host processor.” Instead, Madukkarumukumana paragraph [0028] says that the size of a packet region may be equal to the size of the underlying transport’s MTU. In this case, each I/O reply, consisting of a series of packets, would be put into a corresponding series of MTU frames, so there would be no “bunching” of the I/O replies into the MTU frames.

In view of the above, it is respectfully submitted that the application is in condition for allowance. Reconsideration and early allowance are earnestly solicited.

Respectfully submitted,

/ *Richard C. Auchterlonie* /

Richard C. Auchterlonie, Reg. No. 30,607
NOVAK DRUCE & QUIGG, LLP
1000 Louisiana, 53rd Floor
Houston, TX 77002
713-571-3460 (Telephone)
713-456-2836 (Telefax)
Richard.Auchterlonie@novakdruce.com